arXiv:1608.05368v1 [cs.PL] 17 Aug 2016

Scaling Bounded Model Checking By Transforming Programs With Arrays

Anushri Jana[1], Uday P. Khedker[2], Advaita Datar[1], R. Venkatesh[1], and Niyas C[1]

[1] Tata Research Development and Design Centre, Pune, India
{anushri.jana,advaita.datar,r.venky,niyas.c}@tcs.com
[2] Indian Institute of Technology Bombay, India.
uday@cse.iitb.ac.in

Abstract. Bounded Model Checking is one the most successful techniques for finding bugs in program. However, for programs with loops iterating over large-sized arrays, bounded model checkers often exceed the limit of resources available to them. We present a transformation that enables bounded model checkers to verify a certain class of array properties. Our technique transforms an array-manipulating program in ANSI-C to an array-free and loop-free program. The transformed program can efficiently be verified by an off-the-shelf bounded model checker. Though the transformed program is, in general, an abstraction of the original program, we formally characterize the properties for which the transformation is precise. We demonstrate the applicability and usefulness of our technique on both industry code as well as academic benchmarks.

Keywords: Program Transformation, Bounded Model Checking, Array, Verification.

1 Introduction

Bounded Model Checking is one of the most successful techniques for finding bugs [11] as evidenced by success achieved by tools implementing this technique in verification competitions [2,1]. Given a program P and a property φ, Bounded Model Checkers (BMCs) unroll the loops in P a fixed number of times and search for violations to φ in the unrolled program. However, for programs with loops of large or unknown bounds, bounded model checking instances often exceed the limits of resources available. In our experience, programs manipulating large-sized arrays invariably have such loops iterating over indices of the array. Consequently, BMCs routinely face the issue of scalability in proving properties on arrays. The situation is not different even when the property is an *array invariant* i.e., it holds for every element of the array, a characteristic which can potentially be exploited for efficient bounded model checking.

Consider the example in Figure 1 manipulating an array of structures *a*. The structure has two fields, *p* and *q*, whose values are assigned in the first *for* loop (lines 8–13) such that $a[i].q$ is the square of $a[i].p$ for every index *i*. The second *for* loop (lines 14–17) asserts that this property indeed holds for each element in *a*. This is a *safe* program i.e., none of the assertions admit a counterexample. CBMC [9], a bounded model checker for C, in an attempt to unwind first loop 100000 times, runs out of memory before it

```
1. struct S {
2.     unsigned int p;
3.     unsigned int q;
4. } a[100000];
5. int i,k;

06. main()
07. {
08.  for(i=0; i<100000; i++)
09.  {
10.   k = i;
11.   a[i].p = k;
12.   a[i].q = k * k ;
13.  }

14.  for (i=0; i<100000; i++)
15.  {
16.    assert(a[i].q ==
           a[i].p * a[i].p);
17.  }
18. }
```

Fig. 1: Motivating Example

```
1. struct S{
2.     unsigned int p;
3.     unsigned int q;
4. }x_a;
5. int i_a;
6. int i,k;
7. main()
8. {
9.    i_a = nd(0,99999);

//first loop body
10.  k = nd(0,100000);
11.  i = i_a;
12.  k = i;
13.  (i == i_a)? x_a.p = k : k;
14.  (i == i_a)? x_a.q = k * k : k*k ;
15.  k = nd(0,100000);

//second loop body
16.  i = i_a;
17.   assert(((i==i_a)?x_a.q:nd())
        ==((i==i_a)?x_a.p:nd())
        *((i==i_a)?x_a.p:nd()));
18. }
```

Fig. 2: Transformed Code

even reaches the loop with assertion. In fact, we tried this example with several other model checkers[3] and none of them were able to prove this property because of large loop bounds.

One of the ways of proving this example safe is to show that the property holds for any arbitrary element of the array, say at index i_c. This allows us to get rid of those parts of the program that do not update $a[i_c]$ which, in turn, eliminates the loop iterating over all the array indices. This enables CBMC to verify the assertion without getting stuck in the loop unrollings. Moreover, since i_c is chosen nondeterministically from the indices of a, the property holds for every array element without loss of generality.

This paper presents the transformation sketched above with the aim that the transformed program is easier for a BMC to verify as compared to the original program. The transformation is over-approximative i.e., it give more values than that by the original program. This ensures that if the original program is safe with respect to the chosen property, so is the transformed program. However, the over-approximation raises two important questions spanning practical and intellectual considerations:

1) *Is the proposed approach practically useful? Does the transformation enable a BMC to verify real-world programs, and even academic benchmarks, fairly often?* We provide an answer to this through an extensive experimental evaluation over industry code as well as examples in the array category of SV-COMP 2016 benchmarks. In all the cases, we show that our approach helps CBMC to scale. We further demonstrate the applicability of our technique to successfully identify a large number of false warnings (on an average 73%) reported by a static analyzer on arrays in large programs.

[3] Result for *motivatingExample.c* at https://sites.google.com/site/datastructureabstraction/

2) *Is it possible to characterize a class of properties for which it is precise?*
In order to address this we provide a formal characterization of properties for which the transformation is precise i.e., we state criteria under which the transformed program is unsafe only when the original program is unsafe (Section 6).

To summarize, this paper makes the following contributions:

- A new technique using the concept of *witness index* that enables BMCs to verify array invariant properties in programs with loops iterating over large-sized arrays.
- A formal characterization of properties for which the technique is precise.
- A transformation engine implementing the technique.
- An extensive experimental evaluation showing the applicability of our technique to real-world code as well as to academic benchmarks.

The rest of the paper starts with an informal description of the transformation (Section 2) before we define the semantics (Section 3) and formally state the transformations rules (Section 4). Section 5 and 6, resp., describe the soundness and precision of our approach. Section 7 presents the experimental setup and results. We discuss the related work in Section 8 before concluding in Section 9.

2 Informal Description

Given a program P containing loops iterating over an array a, we transform it to a program P' that has a pair $\langle x_a, i_a \rangle$ of a *witness variable* and a *witness index* for the array and the index such that x_a represents the element $a[i_a]$ of the original program. Further, loops are replaced by their customized bodies that operate only on x_a instead of all elements of a.

To understand the intuition behind our transformation, consider a trace t of P ending on the assertion A_n. Consider the last occurrence of a statement $s : a[e_1] = e_2$ in t. We wish to transform P such that there exists a trace t' of P' ending on A_n with value of i_a equal to that of e_1 and value of x_a equal to that of e_2. We achieve this by transforming the program such that:

- i_a gets a non-deterministic value at the start of the program (this facilitates arbitrary choice of array element $a[i_a]$).
- array writes and reads for $a[i_a]$ gets replaced with witness variable x_a.
- array writes other than $a[i_a]$ gets eliminated and reads gets replaced with non-deterministic value.
- loop body is executed only once either non-deterministically or unconditionally based on loop characteristics. During the execution of the loop body,
  - the loop iterator variable gets the value of i_a or a non-deterministic value (depending on loop characteristics), and
  - all other scalar variables whose values may be different in different iterations gets non-deterministic values.

Figure 2 shows the transformed program P' for the program P of Figure 1. Function `nd(1,u)` returns a non-deterministic value in the range $[l..u]$. In P', the witness index i_a for array a is globally assigned a non-deterministic value within the range of array size (at line 9). In a run of BMC, the assertion is checked for this non-deterministically chosen element $a[i_a]$. To ensure that values for the same index $a[i_a]$ are written and read, we replace array accesses by the witness variable x_a only when the value of index i matches with i_a (lines 13, 14 and 17). We remove loop header but retain loop body. To over-approximate the effect of removal of loop iterations we add non-deterministic assignments to all variables modified in the loop body, at the start of the transformed loop body and also after the transformed loop body (lines 11 and 15). Note that we retain the original assignment statements too (line 12). Since the loops at line 8 and line 14 in the original program iterate over the entire array, we equate loop iterator variable i to i_a (line 11 and 16) and the transformed loop bodies (lines 10–14 and lines 16–17) are executed unconditionally.

We explain the transformation rules formally in Section 4. The transformed program can be verified by an off-the-shelf BMC. Note that each index will be considered in some run of the BMC since i_a is chosen non-deterministically. Hence, if an assertion fails for any index in the original program, it fails in the transformed program too.

3 Semantics

In this section we formalize our technique by explaining the language and defining representation of states.

3.1 Language

We formulate our analysis over a language modelled on C. For simplicity of exposition we restrict our description to a subset of C which includes C style structures and 1-dimensional arrays. Let $\mathbb{C}$, $\mathbb{V}$, and $\mathbb{E}$ be the sets of values computed by the program, variables appearing in the program, and expressions appearing in the program respectively. A value $c \in \mathbb{C}$ can be an integer, floating-point or boolean value. A variable $v \in \mathbb{V}$ can be a scalar variable, a structure variable, or an array variable. We define our program to have only one array variable denoted as a. However, in practice, we can handle multiple arrays in a program as explained in our technical report [22]. We also define $\mathbb{E}_A \subseteq \mathbb{E}$ as set of array expressions of the form $a\,[E\,]$. A lval L can be an array access expression or a variable. Let $c \in \mathbb{C}$, $x, i \in (\mathbb{V} - \{a\})$. We consider assignment statements, conditional statement, loop statement, and assertion statements defined by the following grammar. We define the grammar of our language using the following non-terminals: Program P consists of statements S which may use lvalues L and expressions E. We assume that programs are type correct as per C typing rules.

$$
\begin{array}{lll}
P \rightarrow & S & \\
S \rightarrow & \textbf{if}\ (E)\ S\ \textbf{else}\ S \mid \textbf{if}\ (E)\ S \mid \textbf{for}\ (i = E;\ E;\ E)\ S \mid & \\
& S;\ S \mid L = E \mid \textbf{assert}(E) & (1) \\
L \rightarrow & a[E] \mid x & \\
E \rightarrow & E \oplus E \mid L \mid c &
\end{array}
$$

In practice, we analyze ANSI-C language programs that includes functions, pointers, composite data-structures, all kinds of definitions, and all control structures except multi-dimensional arrays.

3.2 Representing Program States

We define program states in terms of memory location and the value stored in the memory location. We distinguish between *atomic* variables (such as scalar and structure variables) whose values can be copied atomically to a memory location, from non-atomic variables such as arrays. Since we are considering 1-dimensional arrays, the array elements are atomic locations.

Function $\ell(a[i])$ returns the memory location corresponding to the i^{th} index of array a. The memory of an input program consists of all atomic locations:

$$\mathbb{M} = (\mathbb{V} - \{a\}) \cup \left\{ \ell(a[i]) \;\middle|\; 0 \leq i \leq \mathit{lastof}(a) \right\} \tag{2}$$

The function $\mathit{lastof}(a)$ returns the highest index value for array a.

A *program state* is a map $\sigma : \mathbb{M} \to \mathbb{C}$. $[\![e]\!]_\sigma$ denotes the value of expression e in the program state σ.

We transform a program by creating a pair $\langle i_a, x_a \rangle$ for the array a where i_a is the witness index and x_a is the witness variable. The memory of a transformed program with additional variables is:

$$\mathbb{M}' = (\mathbb{V} - \{a\}) \cup \{x_a\} \cup \{i_a\} \tag{3}$$

For a transformed program, a program state is denoted by σ' and is defined over $\mathbb{M}'$.

We explain the relation between states in original and transformed programs using an example. Let a program P have an array variable a and variable k holding the size of the array a. Let the array contain the values $c_i \in \mathbb{C}$, $0 \leq i < n$, where $n \in \mathbb{C}$ is the value of size of the array. Then, a program state, σ at any program point l can be:

$$\sigma = \left\{ (k,n), \left(\ell(a[0]), c_0\right), \left(\ell(a[1]), c_1\right), \ldots, \left(\ell(a[n-1]), c_{n-1}\right) \right\} \tag{4}$$

In the transformed program P', let x_a and i_a be the witness variable and witness index respectively. Let l' be the program point in P' that corresponds to l in P. Then, all possible states in the transformed program at l' are,

$$\begin{aligned}
&\sigma'_0 = \{(k,n), (i_a,0), (x_a,c_0)\} \\
&\sigma'_1 = \{(k,n), (i_a,1), (x_a,c_1)\} \\
&\ldots \\
&\sigma'_{n-1} = \{(k,n), (i_a,n-1), (x_a,c_{n-1})\}
\end{aligned}$$

We now formally define how a state at a program point in the transformed program represents a state at the corresponding program point in the original program.

Definition 1. *Let σ be a state at a program point in P and let σ' be a state at the corresponding program point in P'. Then, σ' represents σ, denoted as $\sigma' \rightsquigarrow \sigma$ if*

$$\sigma' = \{(i_a,c_1)(x_a,c_2)\} \cup \{(y,c) \mid (y,c) \in \sigma, y \in (\mathbb{V} - \{a\})\} \Rightarrow \left(\ell(a[c_1]), c_2\right) \in \sigma$$

$transform(E) =$

$E \equiv (E_1 \oplus E_2)$	$\Rightarrow$	$emit(transform(E_1) \oplus transform(E_2))$	$(3.E_1)$
$E \in \mathbb{E}_A,\ E \equiv a[E_1]$	$\Rightarrow$	$emit((E_1 == i_a)\ ?\ x_a : nd())$	$(3.E_2)$
otherwise	$\Rightarrow$	$emit(E)$	$(3.E_3)$

$transform(S) =$

$S \equiv (L = E),\ L \equiv a[E_1]$	$\Rightarrow$	$emit\big((E_1 == i_a)\ ?\ x_a = transform(E) : transform(E)\big)$	$(3.S_1)$
$S \equiv (L = E), L \not\equiv a[E_1]$	$\Rightarrow$	$emit(L = transform(E))$	$(3.S_2)$
$S \equiv (\textbf{for}(i = E_1;\ E_2;\ E_3)\ S_1),$ $fullarrayaccess(S),$ $u \in loopdefs(S_1)$	$\Rightarrow$	$emit\ \big(u = nd();\ \ //\forall u \in loopdefs(S_1)$ $i = i_a;$ $transform(S_1);$ $u = nd();\ \ //\forall u \in loopdefs(S_1)$ $\big)$	$(3.S_3)$
$S \equiv (\textbf{for}(i = E_1;\ E_2;\ E_3)\ S_1),$ $\neg fullarrayaccess(S),$ $u \in loopdefs(S_1)$	$\Rightarrow$	$emit\ \big(\textbf{if}(nd(0,1))$ $\{\ u = nd();\ \ //\forall u \in loopdefs(S_1)$ $i = nd(loopbound(S));$ $transform(S_1);$ $\}$ $u = nd();\ \ //\forall u \in loopdefs(S_1)$ $\big)$	$(3.S_4)$
$S \equiv (\textbf{if}(E)\ S_1\ \textbf{else}\ S_2)$	$\Rightarrow$	$emit\big(\textbf{if}(transform(E))$ $transform(S_1)\ \textbf{else}\ transform(S_2)\big)$	$(3.S_5)$
$S \equiv (\textbf{if}(E)\ S_1)$	$\Rightarrow$	$emit(\textbf{if}(transform(E))\ transform(S_1))$	$(3.S_6)$
$S \equiv (S_1; S_2)$	$\Rightarrow$	$emit(transform(S_1); transform(S_2))$	$(3.S_7)$
$S \equiv (\textbf{assert}(E))$	$\Rightarrow$	$emit(\textbf{assert}(transform(E)))$	$(3.S_8)$
otherwise	$\Rightarrow$	$emit(S)$	$(3.S_9)$

$transform(P) =$

$P \equiv S$	$\Rightarrow$	$emit\ \big(i_a = nd\,(lastof(a));$ $transform(S)$ $\big)$	$(3.P)$

Fig. 3: Program transformation rules. Non-terminals P, S, E, L represent the code fragment in the input program derivable from them.

Let A_n be the assertion at line n in program P. Let σ be a state reaching A_n in the original program with pair $(\ell\,(a\,[[\![e_1]\!]_\sigma]), [\![e_2]\!]_\sigma)$. Let σ' be the state in transformed program, σ' represents σ. Thus, σ' has two pairs, $(i_a, [\![e_3]\!]_{\sigma'})$ and $(x_a, [\![e_4]\!]_{\sigma'})$ such that $[\![e_3]\!]_{\sigma'} = [\![e_1]\!]_\sigma$ and $[\![e_4]\!]_{\sigma'} = [\![e_2]\!]_\sigma$. Hence, if the assertion A_n holds in transformed program it holds in the original program too.

4 Transformation

The transformation rules are given in Figure 3. A transformed program satisfies the following grammar derived from that of the original program (grammar 1). Let $x, x_a, i_a \in \mathbb{V}$ denote scalar variable, witness variable, and witness index, respectively. Let $c, l, u \in \mathbb{C}$ be values. Then,

$$\begin{aligned} P &\rightarrow\ I\,;\ S \\ I &\rightarrow\ i_a = nd(l,u) \\ S &\rightarrow\ \textbf{if}\ (E)\ S\ \textbf{else}\ S \mid \textbf{if}\ (E)\ S \mid S\,;\ S \mid L = E \mid \textbf{assert}(E) \\ L &\rightarrow\ x \mid x_a \mid i_a \\ E &\rightarrow\ E \oplus E \mid L \mid c \mid nd() \mid nd(l,u) \end{aligned} \tag{5}$$

The non-terminal I represents the initialization statements for witness index. Witness variable is initialized in the scope same as that in the original program.

We use the functions described below in the transformation rules.

- Function *nd* returns a non-deterministically chosen value from the given range l, u; l and u being the lower and upper limit respectively. When range is not provided, *nd* returns a non-deterministic value based on the type of L.
- Function *transform* takes the text derived from a non-terminal and transforms it. Function *emit* shows the actual code that would be emitted. We ignore the details of number of parameters and the type of the parameters of *emit*. We assume that it takes the code emitted by *transform* and possibly some additional statements and outputs the combined code. It has been used only to distinguish the transformation time activity and run time activity. For example, the boolean conditions in cases $3.E_2$ and $3.S_1$ are not evaluated by the body of function *transform* but is a part of the transformed code and is evaluated at run time when the transformed program is executed. Similar remarks apply to the **if** statements and other operations inside the parenthesis of *emit* function.
- Function $fullarrayaccess(S)$ analyzes[4] the characteristics of the loop S.
  - When the loop S accesses array a completely, $fullarrayaccess(S)$ returns true. This means that loop either reads or write all the indices of the array.
  - When the loop S accesses array a partially, $fullarrayaccess(S)$ returns false. This means that the loop may not access all the indices or some indices are being read while some other indices are being written.
  - When loop S do not access an array, $fullarrayaccess(S)$ returns false.
- Function $loopdefs(S)$ returns the over-approximated set of variables modified in the loop S.
  - Scalar variables are included in this set if they appear on the left hand side of any assignment statement in S (except when the RHS is a constant).
  - Loop iterator variable i of loop S is not included in this set.
  - Array variable a is included in this set when the array access expression appears on the left hand side of an assignment and the value of index expression is different from the current value of the loop iterator i.

[4] Analysis can be over-approximated.

– Function $lastof(a)$ returns the highest index value for array a.

With the above functions, the transformation rules are easy to understand. Here we explain non-trivial transformations.

– To choose an array index for a run, witness index (i_a) is initialized at the start of the program to a non-deterministically chosen value from the range of the indices of the array (case 3.P). This value determines the array element ($a[i_a]$) represented by the witness variable (x_a).
– An array access expression in LHS or RHS is replaced by the witness variable (x_a) provided the values of the witness index and index expression of the array access expression match. If the values do not match, it implies that the element accessed is not at the non-deterministically chosen index i_a. Hence for any other index the assignment does not happen (case 3.S_1). Similarly, when any other index is read in RHS, it is replaced with a non-deterministic value (case 3.E_2).
– Loop iterations are eliminated by removing the loop header containing initialization, test, and increment expression for loop iterator variable. The loop bodies are transformed as follows :
  - Each variable in the set returned by $loopdefs(S)$ is assigned a non-deterministic value at the start of the loop body and also after the loop body. These assignments ensure that values dependent on loop iterations are over-approximated when used inside or outside the loop body.
  - The loop iterator i is a special scalar variable. A loop S where $fullarrayaccess(S)$ holds (case 3.S_3) essentially means that loop bound is same as the array size and array is accessed using loop iterator as index. Hence it is safe to replace array access with x_a where the values of loop iterator and index expression match. To ensure this we equate loop iterator with i_a. This models the behaviour of the original program precisely. However, when $fullarrayaccess(S)$ does not hold (case 3.S_4), we assign loop iterator i to a non-deterministically chosen value from the loop bound.
  - Each statement in the loop body is transformed as per the transformation rules.
  - Finally, the entire loop body is made conditional using a non-deterministically chosen true/false value when $fullarrayaccess(S)$ does not hold. This models the partial accesses of array indices which imply that some of the values defined before the loop may reach after the loop. However, the transformed loop body is unconditionally executed when $fullarrayaccess(S)$ holds.

5 Soundness

This section outlines the claim that the proposed transformation is sound, i.e. if the transformed program is safe, then so is the original program. As discussed in Section 3, the soundness is immediate if the abstract states "represent" the original states. We, therefore, prove that the proposed transformations ensure that the *represents* relation, $\rightsquigarrow$, holds between abstract and original states. For the base case, we prove that $\rightsquigarrow$ holds in the beginning - before applying any transformation (Lemma 1). In the inductive step, we prove that if $\rightsquigarrow$ holds at some stage during the transformation, then the subsequent transformation continues to preserve $\rightsquigarrow$ (Lemma 3). We prove this by structural induc-

tion on program transformations. We prove that each transformed expression is over-approximated when $\rightsquigarrow$ holds in (Lemma 2). Detailed proof is provided in our technical report [22].

Lemma 1. *Let the start of the original program (i.e. the program point just before the code derivable from non-terminal S in production $P \rightarrow S$ in grammar defined in equation(1) be denoted by l. The corresponding program point in the transformed program P', denoted by l', is just after I and just before the non-terminal S in production $P \rightarrow I\,;\ S$ (Grammar in equation 5). Let σ and σ' be the states at l and l' in P and P' respectively. Then, $\sigma'_{l'} \rightsquigarrow \sigma_l$.*

Proof Outline. Since the initial values of non array variables are preserved, the initial value of the element of array $a[i_a]$ is assigned to x_a, and i_a is non-deterministically chosen, the lemma holds.

Lemma 2. *Let σ_l be a state at a program point l in P and $\sigma'_{l'}$ be a state at the corresponding program point l' in transformed program P'. Consider an arbitrary expression $e \in \mathbb{E}$ just after l in original program P. Then,*

$$\sigma'_{l'} \rightsquigarrow \sigma_l \Rightarrow [\![\mathit{transform}(e)]\!]_{\sigma'_{l'}} \supseteq [\![e]\!]_{\sigma_l}.$$

Proof Outline. Since e is derived from E (grammar 1), the over-approximation of values can be proved by structural induction on the productions for E.

Lemma 3. *Let l and m be the program points just before and after a statement s in P and let σ_l and σ_m be the states at l and m respectively. Let l' and m' be the program points just before and after the corresponding transformed statement $\mathit{transform}(s)$ in P'. Let $\sigma'_{l'}$ and $\sigma'_{m'}$ be the states at l' and m' respectively. Then, $\sigma'_{l'} \rightsquigarrow \sigma_l \Rightarrow \sigma'_{m'} \rightsquigarrow \sigma_m$.*

Proof Outline. Since statement s is derived from non-terminal S in the grammar 1 the lemma can be proved by structural induction on S.

Theorem 1. *If the assertion A_n is violated in the original program P, then it will be violated in transformed program P' also.*

Proof. Let the assert get violated for some $a[c]$. Since i_a is initialized non-deterministically it can take the value c and we have shown in Lemma 2 that all expressions in P' are over-approximated. Lemma 1 and Lemma 3 ensure the premise for Lemma 2. Hence the theorem follows.

6 Precision

We characterize the assertions for which our transformation is precise – an assertion will fail in P' if and only if it does so in P. We denote such an assertion as A_n^{inv}. We focus on A_n^{inv} in a loop. A program can have array accesses outside loops too. In such cases we do not claim precision; as per our experience such situations are rare in programs with large-sized arrays.

Our transformations replace array access expressions and loop statements while the statements involving scalars alone outside the loop remain unmodified. Hence precision criteria need to focus on the statements within loops and not outside it.

Let assertion A_n^{inv} be in loop statement $S_{A_n^{inv}}$. Let $\mathbb{V}_{\mathsf{imp}}$ be the set of variables and $\mathbb{E}_{\mathsf{imp}}$ be the set of array access expressions on which A_n^{inv} is data or control dependent within the loop $S_{A_n^{inv}}$. Let the set of loop statements from where definitions reach A_n^{inv} be denoted by $\mathbb{S}_{\mathsf{def}}$, note that this set is a transitive closure for data dependence. Our technique is precise when:

- $\mathit{fullarrayaccess}(S)$ holds for each $S \in \{S_{A_n^{inv}}\} \cup \mathbb{S}_{\mathsf{def}}$ (rule l_1)
- If $a[e] \in \mathbb{E}_{\mathsf{imp}}$ then
  - the index expression $e = i$ where i is the loop iterator of loop $S_{A_n^{inv}}$ (rule a_2)
  - $a \notin \mathit{loopdefs}(S)$ where $S \in \{S_{A_n^{inv}}\} \cup \mathbb{S}_{\mathsf{def}}$ (rule a_3)
- If $x \in \mathbb{V}_{\mathsf{imp}}$ then $x \notin \mathit{loopdefs}(S)$ where $S \in \{S_{A_n^{inv}}\} \cup \mathbb{S}_{\mathsf{def}}$ (rule s_4)
- For an assignment statement of the form $a[e_1] = e_2$ in loop S where $S \in \mathbb{S}_{\mathsf{def}}$,
  - if e_2 is an array access expression then it must be of the form $a[i]$ where i is the loop iterator of loop S (rule d_5)
  - if e_2 is x then $x \notin \mathit{loopdefs}(S)$ where $S \in \mathbb{S}_{\mathsf{def}}$ (rule d_6)

Theorem 2. *If the assertion A_n^{inv}; that satisfies above rules; holds in the original program P, then it will hold in the transformed program P' also.*

Proof. The transformed program is over-approximative because our transformation rules ($3.S_3$, $3.S_4$, $3.E_2$) introduce non-deterministic values. We prove this theorem by showing that if assertion is of the form A_n^{inv} then none of these transformation rules introduce non-deterministic values in the transformed program.

- Since rule l_1 holds unconditionally, case $3.S_4$ will not apply. Hence no extra paths are added in transformed program. Also, since case $3.S_3$ applies, assignment $i = i_a$ will be added for $S_{A_n^{inv}}$ and the loop statements in $\mathbb{S}_{\mathsf{def}}$.
- When rule a_2 holds, since rule l_1 holds $a[e]$ get replaced by x_a always (case $3.E_2$).
- When rule a_3 holds, assignment $x_a = nd()$ is not added (case $3.S_3$).
- When rule s_4 holds, assignment $x = nd()$ is not added (case $3.S_3$).
- When rule d_5 holds, since rule l_1 holds $a[e]$ in RHS gets replaced with x_a (case $3.E_2$).
- When rule d_6 holds, scalars in RHS are not assigned with a non-deterministic value.

Note that rule s_4 is a very strong condition to ensure that non-deterministic values do not reach A_n^{inv}. We can relax this rule when $x \in \mathit{loopdefs}(S_{A_n^{inv}})$ under these two conditions:

- definition of x appears before the assert statement in the loop
- x is defined with a constant or using loop iterator i only.

None of the transformation rules replace variable x. Definition of x to non-deterministic value ($x = nd()$) gets re-defined by original assignment (retained in the transformed loop body) appearing before the assert statement. Since x is defined with a constant or i ($i = i_a$ is added for $S_{A_n^{inv}}$), its value is not over-approximated.

Table 1: Results on SV-COMP Benchmark Programs.

#programs = 118	#correct true	#correct false	#incorrect true	#incorrect false	#no result
Expected Results	84	34	-	-	0
$CBMC_{\alpha}$	47	6	6	0	59
$CBMC_{\beta}$	9	5	0	0	104
Transformation+$CBMC_{\beta}$	25	34	0	59	0
$CBMC_{\alpha}$ - SV-COMP2016 (unsound) CBMC, $CBMC_{\beta}$ - sound CBMC 5.4					

7 Experimental Evaluation

We have implemented our transformation engine using static analysis[5]. It supports ANSI-C programs with 1-dimensional arrays. The experiments are performed on a 64-bit Linux machine with 16 Intel Xeon processors running at 2.4GHz, and 20GB of RAM. More details of optimization and implementation, including handling of multiple arrays, are provided in our technical report [22].

Our transformation engine outputs C programs. Although we could take any off-the-shelf BMC for C program to verify the transformed code, we use CBMC in our experiments as it is known to handle all the constructs of ANSI-C. We discuss the results of our experiments on academic benchmarks and industry codes. For want of space, we omit the results of various BMCs on patterns from industry code; those results are shared in our technical report [22].

7.1 Experiment 1 : SV-COMP Benchmarks

SV-COMP benchmarks [27] contain an established set of programs under various categories intended for comparing software verifiers. Results for *ArraysReach*[6] from the *array* category for CBMC used in SV-COMP 2016 ($CBMC_{\alpha}$), CBMC 5.4 ($CBMC_{\beta}$) and CBMC 5.4 on transformed programs (Transformation+$CBMC_{\beta}$) are consolidated[7] in Table 1. *ArraysReach* has 118 programs. $CBMC_{\alpha}$, an unsound version of CBMC, gave correct results for 53 programs. However, $CBMC_{\beta}$ gave correct results for 14 programs. We compare the results of Transformation+$CBMC_{\beta}$ on three criteria:

- Scalability: it scaled up for all 118 programs.
- Soundness: it gave sound results for all 118 programs. For the 6 program for which $CBMC_{\alpha}$ gave unsound results, our results are not only sound but are also precise.
- Precision: it gave precise results for 59 programs. Out of these $CBMC_{\alpha}$ ran out of memory for 45 programs ($CBMC_{\alpha}$ ran out of memory for 14 additional programs). On the other hand, 22 true programs reported correctly by $CBMC_{\alpha}$ were verified as false by Transformation+$CBMC_{\beta}$. Transformation+$CBMC_{\beta}$ verified 25 program as true which did not include 8 of programs reported correctly as true by $CBMC_{\beta}$.

[5] PRISM, a static analyzer generator developed at TRDDC, Pune [10,23]

[6] Programs in *ArrayMemSafety* access arrays without using index and cannot be transformed.

[7] Case by case results available at https://sites.google.com/site/datastructureabstraction/home/sv-comp-benchmark-evaluation-1

Table 2: Real-life Application Evaluation

Application details				Sliced+CBMC			Sliced +Transformation +CBMC			% False Positive Reduction
Name	Size (LoC)	$\%loop^{full}$	#Asserts	#P	#F	#T	#P	#F	#T	
navi1	1.54M	100	63	0	0	63	52	1	10	82.5
navi2	3.3M	93.4	103	0	0	103	95	1	7	92.2
icecast_2.3.1	336K	59.1	114	0	0	114	53	61	0	46.5

$loop^{full}$ - loop S where $fullarrayaccess(S)$ holds,
P - Assertion Proved, F - Assertion Failed, T - Timeout

Our technique is imprecise for the other 59 of 118 programs as they do not comply with the characterization of precision provided in Section 6. As can be seen, there is a trade-off between scalability and precision. From the view point of reliability of results, soundness is the most desirable property of a verifier. Our technique satisfies this requirement. Further, it not only scales up but is also precise implying its practical usefulness.

7.2 Experiment 2 : Real-life Applications

We applied our technique on 3 real-life applications - *navi1* and *navi2* are industry codes implementing the navigation system of an automobile and icecast_2.3.1 is an open source project for streaming media [21]. We appended assertions using *null pointer dereference* (NPD) warnings from a sound static analysis[8] tool as follows. Lets say the dereference expression is $*a[i].p$. A statement $assert(a[i].p != null)$ is added in the code just before statement containing dereference expression.

We ran CBMC on these applications with a time out of 30 minutes. CBMC did not scale on the original as well as the sliced programs. We ran our transformation engine on sliced programs. Table 2 shows the consolidated results of our experiments. Out of 280 assertions, sliced+transformation+CBMC proved 200 assertions taking 12 minutes on average for transformation+verification. This is a much less in comparison to the time given to CBMC for sliced programs (sliced+CBMC), which was 30 minutes.

To verify the correctness of our implementation, we analyzed the warnings manually. We found that all 280 warnings were false, implying that all the assertions should have been proved successfully.

– CBMC could scale up for such large applications because there are no loops in transformed programs. However, CBMC could not scale for 17 cases even after transformation because of the presence of a long recursive call chain of calls through function pointers.

[8] TCS Embedded Code Analyzer (TCS ECA)
http://www.tcs.com/offerings/engineering_services/Pages/TCS-Embedded-Code-Analyzer.aspx

– CBMC could not prove 63 of the assertions since array definitions reaching at the assertion were from the loops where *fullarrayaccess*(S) did not hold. Hence the witness variable takes over-approximated values.
– CBMC proved 200 assertions, where all the conditions for precision mentioned in Section 6 get fulfilled. In these experiments, we checked for the NPD property which is value-independent. Moreover, we found that the assertions inserted by us are not control-dependent on any scalar.

Note that the number of false warnings eliminated in an application is proportional to the number of loops for which *fullarrayaccess* (S) hold. Over a diverse set of applications, we found that our technique could eliminate 40-90% of false warnings. This is a significant value addition to static analysis tools that try to find defects and end up generating a large number of warnings. In fact, our own effort grew out of the need of handling warnings that were generated by our proprietary static analysis tool, a large fraction of which were false positives.

8 Related Work

The literature on automated reasoning about array-manipulating code can be broadly categorized into *analysis* and *verification*. Most methods that analyze programs manipulating arrays [8,18,20,13,25] are based on abstract interpretation. Cornish et al. [12] transform a program to remove arrays and discover non-trivial universally quantified loop invariants by analyzing the transformed program using off-the-shelf abstract scalar analysis. Since they create additional blocks for each value of *summary variable*, the program size increases considerably raising concerns about scalability. Similar to our approach, Monniaux et al. [26] transform array programs by replacing array operations with a scalar. However they keep loops. These programs are then analyzed using methods producing invariants (*back-ends*). CBMC did not scale up on the transformed "array copy" example (10000 loop bound) given in the paper, suggesting that scalability is a concern with this technique too. However, using our technique CBMC scaled for the same program.

Dillig et al. [15] introduced fluid updates of arrays in order to do away with strong and weak updates. Their technique uses indexed locations along with bracketing constraints, a pair of over- and under-approximative constraints, to specify the concrete elements being updated. In another work [16], they propose an automatic technique to reason about contents of arrays (or containers, in general). However, they introduce an abstraction to encode all values that (a subset of) elements may have. In contrast, since our technique choses only one representative element to work with, we can capture its value precisely.

Template-base methods [7,19] have been very useful in synthesizing invariants but these techniques are ultimately limited by a large space of possible templates that must be searched to get a good candidate template. This has also led to semi-automatic approaches, such as [17], where the predicates are usually suggested by the user. Our approach, however, is fully automatic and proves safety by solving a bounded model checking instance instead of computing an invariant explicitly.

Verification tools based on CEGAR have been applied successfully to certain classes of programs, e.g., device drivers [6]. However, this technique is orthogonal to ours. In fact, a refinement framework in addition to our abstraction would make our technique complete. Several other techniques have been used to scale BMCs to tackle complex, real-world programs such as acceleration [24] and loop-abstraction [14]. But these techniques are not shown to be beneficial in abstracting complex data structures. Booster [4], a recent tool for verifying C-like programs handling arrays, integrates acceleration and lazy abstraction with interpolants for arrays [5,3]. It exploits acceleration techniques to compute an exact set of reachable states, whenever possible, for programs with arrays. For instance, their technique works on $simple^0_{\mathcal{A}}$ programs [5]. However, there are syntactic restrictions that limit the applicability of acceleration in general for programs handling arrays. Note that Booster uses acceleration, instead of abstraction-based procedures, for want of a precise solution (not involving over-approximations). Since our technique is also precise for a characterizable class of programs, it is certainly possible to gainfully combine the two techniques in order to handle a larger class of programs than what either of them can handle in isolation.

9 Conclusions and Future Work

Verification of programs with loops iterating over arrays is a challenging problem because of large sizes of arrays. We have explored a middle ground between the two extremes of relying completely on dynamic approaches of using model checkers on the one hand and using completely static analysis involving complex domains and fix point computations on the other hand. Our experience shows that using static analysis to transform the program and letting the model checkers do the rest is a sweet spot that enables verification of properties of arrays using an automatic technique that is generic, sound, scalable, and reasonably precise.

Our experiments show that the effectiveness of our technique depends on the characteristics of programs and properties sought to be verified. We are able to eliminate 40-90% of false warnings from diverse applications. This is a significant value addition to static analysis that try to find defects and end up generating a large number of warnings which need to be resolved manually for safety critical applications. Our effort grew out of our own experience of such manual reviews which showed a large number of warnings to be false positives.

We plan to make our technique more precise by augmenting it with a refinement step to verify the programs that are reported as unsafe by our current technique. Finally, we wish to extend our technique on other data structures such as maps or lists.

References

1. 2015 4th International Competition on Software Verification. `http://sv-comp.sosy-lab.org/2015/results/`. [Online; accessed 25-Jun-2016].
2. 2016 5th International Competition on Software Verification. `http://sv-comp.sosy-lab.org/2016/results/results-verified/`. [Online; accessed 25-Jun-2016].

3. F. Alberti, R. Bruttomesso, S. Ghilardi, S. Ranise, and N. Sharygina. Lazy abstraction with interpolants for arrays. In *LPAR*, 2012.
4. F. Alberti, S. Ghilardi, and N. Sharygina. Booster: An acceleration-based verification framework for array programs. In *ATVA*. 2014.
5. F. Alberti, S. Ghilardi, and N. Sharygina. Decision procedures for flat array properties. In *TACAS*. 2014.
6. T. Ball and S. K. Rajamani. The slam project: debugging system software via static analysis. In *ACM SIGPLAN Notices*, volume 37, 2002.
7. D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Invariant synthesis for combined theories. In *VMCAI*, 2007.
8. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In *The Essence of Computation*. 2002.
9. CBMC. `http://www.cprover.org/cbmc/`. [Online; accessed 25-Jun-2016].
10. B. Chimdyalwar and S. Kumar. Effective false positive filtering for evolving software. In *ISEC*, 2011.
11. F. Copty, L. Fix, R. Fraer, E. Giunchiglia, G. Kamhi, A. Tacchella, and M. Y. Vardi. Benefits of bounded model checking at an industrial setting. In *CAV*, 2001.
12. J. R. M. Cornish, G. Gange, J. A. Navas, P. Schachte, H. Søndergaard, and P. J. Stuckey. Analyzing array manipulating programs by program transformation. In *LOPSTR*. 2014.
13. P. Cousot, R. Cousot, and F. Logozzo. A parametric segmentation functor for fully automatic and scalable array content analysis. In *ACM SIGPLAN Notices*, volume 46, 2011.
14. P. Darke, B. Chimdyalwar, R. Venkatesh, U. Shrotri, and R. Metta. Over-approximating loops to prove properties using bounded model checking. In *DATE*, 2015.
15. I. Dillig, T. Dillig, and A. Aiken. Fluid updates: Beyond strong vs. weak updates. In *Programming Languages and Systems*. 2010.
16. I. Dillig, T. Dillig, and A. Aiken. Precise reasoning for programs using containers. *ACM SIGPLAN Notices*, 46(1), 2011.
17. C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In *ACM SIGPLAN Notices*, volume 37, 2002.
18. D. Gopan, T. Reps, and M. Sagiv. A framework for numeric analysis of array operations. *ACM SIGPLAN Notices*, 40(1), 2005.
19. S. Gulwani, B. McCloskey, and A. Tiwari. Lifting abstract interpreters to quantified logical domains. In *POPL*, 2008.
20. N. Halbwachs and M. Péron. Discovering properties about arrays in simple programs. In *ACM SIGPLAN Notices*, volume 43, 2008.
21. Ice Cast. `http://icecast.org/`. [Online; accessed 25-Jun-2016].
22. A. Jana, U. P. Khedker, A. Datar, R. Venkatesh, and N. C. Scaling bounded model checking by transforming programs with arrays. *CoRR*, arXiv:1606.06974, 2016.
23. S. Khare, S. Saraswat, and S. Kumar. Static program analysis of large embedded code base: an experience. In *ISEC*, 2011.
24. D. Kroening, M. Lewis, and G. Weissenbacher. Under-approximating loops in c programs for fast counterexample detection. In *CAV*, 2013.
25. J. Liu and X. Rival. Abstraction of arrays based on non contiguous partitions. In *VMCAI*, 2014.
26. D. Monniaux and F. Alberti. A simple abstraction of arrays and maps by program translation. In *Static Analysis*. 2015.
27. SV-COMP 2016 Benchmarks. `https://sv-comp.sosy-lab.org/2016/benchmarks.php`. [Online; accessed 25-Jun-2016].